

PCT

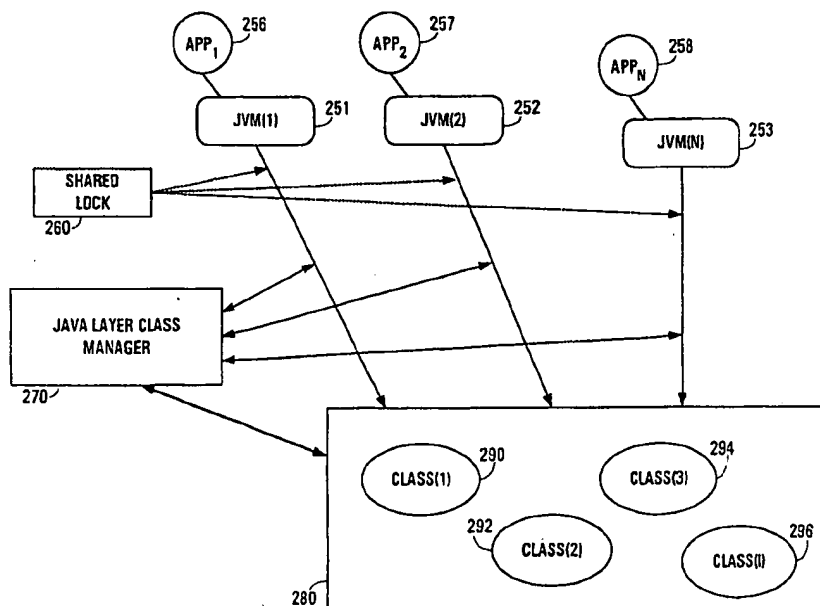
WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : G06F 9/445	A1	(11) International Publication Number: WO 00/52572 (43) International Publication Date: 8 September 2000 (08.09.00)
(21) International Application Number: PCT/US99/28859 (22) International Filing Date: 6 December 1999 (06.12.99) (30) Priority Data: 09/260,259 1 March 1999 (01.03.99) US (71) Applicant: SONY ELECTRONICS, INC. [US/US]; 1 Sony Drive, Park Ridge, NJ 07656 (US). (72) Inventor: LI, Qiaoyun; 19400 Sorenson Avenue #101, Cupertino, CA 95014 (US). (74) Agents: GALLENSON, Mavis, S. et al.; Ladas & Parry, Suite 2100, 5670 Wilshire Boulevard, Los Angeles, CA 90036-5679 (US).		(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: A COMPUTER IMPLEMENTED METHOD**(57) Abstract**

A computer-implemented method and system for allowing Java classes to be shared among many Java virtual machines (JVMs) including a communication system allowing Java and native applications to readily interoperate. An implementation of the JVM on an operating system platform, e.g., the Aperios AV/OS, allows a variety of applications including desktop applications, applets and Internet based applications, home networking applications, MHEG-6 applets, gaming, gaming applications and next generation audio visual applications to operate with the operating system. The present invention provides a shared memory pool (SMP) into which a JVM and store and register a particular Java class. The stored and registered Java class is then accessible by other JVMs using the SMP and a Java layer class manager that is implemented in software. The Java layer class manager requires other JVMs to access a key for the stored class in order to synchronize access to the Java class among several installed and operating JVMs of the computer system. By sharing common Java classes in this fashion, the memory resource overhead required to operate multiple JVMs on a common computer system is drastically reduced thereby allowing a multiple JVM platform to be operable on an embedded computer system. A novel communication method is also disclosed for communicating information between a JVM application and a native application using the computer system's operating system. The novel communication method also allows multiple JVM applications to communicate using the shared memory pool. These functions are incorporated into a JavaLayer that supports the full PersonalJava™ platform.



Best Available Copy

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

A COMPUTER IMPLEMENTED METHOD

5

BACKGROUND OF THE INVENTIONRELATED US APPLICATION

The present application claims priority of United States Provisional Patent Application serial number 60/108,648, filed November 14, 1998, entitled "JavaLayer: A framework for Implementation of PersonalJava on AperiOS Audio Video Platform (AV/OS)," by Qiaoyun (Liz) Li, and attorney docket number 50MGBG2.

FIELD OF THE INVENTION

15 The present invention relates to the field of software architectures for a computer system running multiple Java Virtual Machines (JVMs). More specifically, the present invention relates to a software architecture that can be used for audio/visual digital devices including embedded systems. Described in this paper in one aspect is a computer-implemented sharing of Java classes for increased memory efficiency and communication method.

20 RELATED ART

Many computer operating systems require Java for various applications including home networking applications, MPEG applets and other desktop and Internet applications. One such home networking application is the HAVi (Home Audio/Video Interoperability) application as described in a reference
25 entitled "HAVi: an IT architecture for interoperability in the home network," by Rodger Lea and Aaron Ludtke, SRF97, which is cited herein as background.

-2-

Given the wide range of target platforms and applications that can be hosted by computer operating systems, in both IT (Information Technology) and consumer electronics markets, it is desired that PersonalJava (pJava) be supported on various types of devices such as, for instance, set-top boxes, consumer
5 electronic (CE) devices and digital televisions (DTVs).

Moreover, the implementation of the standard Java Virtual Machine (JVM) on personal computers (PCs) or on workstations allows applications to be developed in Java by programmers who do not necessarily understand the
10 operating system of the host computer system. In one example, this host computer operating system is the Aperios operating system available from Sony Corporation. Because Aperios' AVIOS is a highly configurable, real-time and micro-kernel platform, it would be advantageous that highly compact and memory efficient implementations of the JVM be implemented on top of it. For
15 instance, it would be advantageous to port PersonalJava to the Aperios's AVIOS environment. Heretofore, this has not been done for a number of reasons.

One problem is that PersonalJava does not currently support Audio.
20 Video (AV) applications, which need more flexible functionality to adequately handle certain functionality such as volume data handling and transfer, device and stream control management, message passing, event management and communication with native applications. Therefore, it would be advantageous to design and build a dedicated environment for the implementation of a Java
25 system infrastructure which would help AV applications.

-3-

There are other problems with respect to running Java systems on dedicated environments, e.g., an embedded computer system or other intelligent electronic device. Namely, embedded computer systems do not have the memory resources required to support Java systems. Figure 1 illustrates a software architecture where a single computer system is running multiple JVMs 11-14 simultaneously. Each JVM installs, into memory, its own copy of the Java classes which are required to allow the respective JVM to operate. For instance, JVM 11 installs classes 21 for its own use, JVM 12 installs classes 22 for its own use, JVM 13 installs classes 23 for its own use and JVM 14 installs classes 24 for its own use. However, many of the installed Java classes are copies of the same class, e.g., class A and class B are copied and are individually used by many of the JVMs. The size of these classes, and the fact that each JVM needs its own copy, makes it very difficult to run multiple JVMs simultaneously on a single computer system. This is especially true for embedded computer systems which lack robust memory resources. What is needed is a system and method for allowing efficient use of memory resources when operating Java on a dedicated environment.

Moreover, when Java systems are running on dedicated environments, there is presently no effective way to pass messages between native applications and JVM applications or between two or more JVM applications. What is needed is a system and method for allowing efficient message passing between native applications and JVM applications, or between two or more JVM applications, within a dedicated environment.

-4-

SUMMARY OF THE INVENTION

Accordingly, the present invention provides a system and method for allowing efficient use of memory resources when operating Java on a dedicated environment, e.g., an embedded computer system or other intelligent electronic device. The present invention further provides a system and method for allowing efficient message passing functionality between native applications and JVM applications, and between two or more JVM applications, within the dedicated environment. Specifically, what is disclosed herein is a novel class sharing model for minimizing the memory footprint of multiple JVMs running on a dedicated computer system, e.g., an embedded system. Embodiments of the present invention also include a model for mapping different physical devices installed with their own file management and a model for external applications to communicate with PersonalJava.

To reduce the memory footprint of the JVMs, a computer-implemented method and system are described for allowing Java classes to be shared among many Java Virtual Machines (JVMs) and this system includes a communication system allowing Java and native applications (or two or more JVM applications) to readily interoperate. Various features of JAVA, like "write once, run everywhere," compactness, standardized APIs (application program interfaces) and the availability of advanced tool kits all make JAVA a very attractive platform for applications. A software framework described herein, referred to as "JavaLayer," (or Java layer) allows various coexisting applications to share resources and provide an optimal run-time environment. The

-5-

JavaLayer enables Java and native applications to interoperate easily, manage class sharing and aids by efficient garbage collection.

An implementation of the JVM on an operating system platform, e.g., the Aperios AV/OS, allows a variety of applications including desktop applications, applets and Internet based applications, home networking applications, MPEG applets, gaming, gaming applications and next generation audio visual applications to operate with the operating system. The present invention provides a shared memory pool (SMP) into which a JVM can store and register a particular Java class for subsequent use by itself or by other JVMs. The stored and registered Java class is accessible by other JVMs using the SMP and a Java layer class manager (JCM) that is implemented in software. The JCM requires other JVMs to access a key for the stored class in order to synchronize access to the Java class among several installed and operating JVMs of the computer system. By sharing common Java classes in this fashion, the memory resource overhead required to operate multiple JVMs on a common computer system is drastically reduced thereby allowing a multiple JVM platform to be operable on an embedded computer system.

A novel communication method is also disclosed for communicating information between a JVM application and a native application, or between two or more JVMs, using the computer system's operating system and the SMP. The novel communication method allows multiple JVM applications to communicate using the shared memory pool. These functions are incorporated into a JavaLayer that supports the full PersonalJava™ platform.

-6-

More specifically, in a computer system having a processor coupled to a bus, a memory device coupled to the bus and a display screen coupled to the bus, embodiments of the present invention include a method of sharing Java classes comprising the steps of: a) a first Java Virtual Machine (JVM) loading an instantiation of a first Java class into a shared memory pool; b) registering a name of the first Java class into a name table for the shared memory pool; c) a second JVM querying the name table to determine if the first Java class is stored in the shared memory pool, wherein the second JVM is operating on the computer system simultaneously with the first JVM; and d) the second JVM accessing and utilizing the instantiation of the first Java class from the shared memory pool in lieu of creating a separate instantiation of the first Java class for its own use.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates a prior art software configuration of Java virtual machines (JVMs) including each JVM having its own separate memory allocation for Java class storage.

5

Figure 2 illustrates a general computer system utilized in accordance with an embodiment of the present invention.

Figure 3 is a logical block diagram of the system architecture of the computer system used by the present invention.

Figure 4 is a logical block diagram of the JavaLayer software model used by the computer system of the present invention.

Figure 5 is a logical block diagram of the multi-JVM JavaLayer in accordance with an embodiment of the present invention including the metaspaces with which it interoperates.

Figure 6 is another logical block diagram of the JavaLayer architecture supporting multiple applications on one JVM simultaneously in accordance with an embodiment of the present inventions.

Figure 7 is a block diagram of the JVMs interfacing with the shared class memory pool of the present invention using the JavaLayer class manager.

Figure 8 illustrates the physical and logical memory organization of the class sharing system of the present invention.

5 Figure 9 illustrates an example of class sharing of the class memory pool across multiple applications using the JavaLayer class manager (JCM) in accordance with the present invention.

10 Figure 10 is a flow diagram illustrating steps of the class sharing process of the present invention using the shared memory pool and the JavaLayer class manager to share classes between multiple JVMs.

15 Figure 11 illustrates a logical diagram of real-time coordination between native applications and a JVM in accordance with the message transfer method of an embodiment of the present invention.

20 Figure 12 illustrates a logical diagram of real-time coordination between two JVMs in accordance with the message transfer method of an embodiment of the present invention.

 Figure 13 illustrates a timing diagram of events of the real-time coordination between two JVMs in accordance with the message transfer method of an embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

In the following detailed description of the present invention, a computer-implemented method and system for sharing Java classes across multiple Java virtual machines on the same computer system and a method for

5 communication between Java and native applications (or between two or more JVMs), numerous specific details are set forth in order to provide a thorough understanding of the present invention. However, it will be recognized by one skilled in the art that the present invention may be practiced without these specific details or with equivalents thereof. In other instances, well known

10 methods, procedures, components, and circuits have not been described in detail as not to unnecessarily obscure aspects of the present invention.

The following documents are referenced as background material and as such are incorporated herein by reference: Aperios™ v.1.1.1 User Guide, Sony

15 1998; "HAVI: an IT architecture for interoperability in the home network," by Rodger Lea and Aaron Ludtke, SRF97; "The Java Virtual Machine Specification," by Tim Lindholm & Frank Yellin, Addison Wesley, 1996; and "With HotSpot, the JVM is faster than a speeding bullet," by Eric Armstrong, JavaWorld, March 25, 1998.

20

Embodiments of the present invention include a dedicated framework, called JavaLayer, that allows various coexisting applications to share resources and provide an optimal run-time environment. The JavaLayer enables Java and native applications to interoperate easily, manages class

25 sharing and aids efficient garbage collection. It also allows two or more JVMs

-10-

to interoperate. In one embodiment, JavaLayer supports the full PersonalJava platform. An implementation of the Java Virtual Machine (JVM) on the Aperios AVIOS platform allows a variety of applications including desktop applications, applets and Internet based applications, home networking applications, MPEG-
5 6 applets, gaming applications and next generation AV oriented applications to run on Aperios. Various features of Java like "write once, run everywhere", compactness, standardized APIs and the availability of advanced tool kits, all make this platform very attractive for applications.

10

NOTATION AND NOMENCLATURE

Some portions of the detailed descriptions which follow are presented in terms of procedures, steps, logic blocks, processing, and other symbolic representations of operations on data bits within a computer memory. These descriptions and representations are the means used by those skilled in the
15 data processing arts to most effectively convey the substance of their work to others skilled in the art. A procedure, computer executed step, logic block, process, etc., is here, and generally, conceived to be a self-consistent sequence of steps or instructions leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not
20 necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated in a computer system. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

25

It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussions, it is appreciated that

5 throughout the present invention, discussions utilizing terms such as "processing" or "computing" or "translating" or "calculating" or "determining" or "scrolling" or "displaying" or "recognizing" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities

10 within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

COMPUTER SYSTEM 112

15 Aspects of the present invention, described below, are discussed in terms of steps executed on a computer system, also called embedded system, (e.g., process 410 of Figure 10) for sharing Java classes. Within the present invention, the computer system can be integrated within a portable electronic device or system, e.g., a personal digital assistant, a portable computer system

20 (e.g., a laptop, a palm sized device), a portable consumer based electronic device such as a mobile phone or car navigator. As an embedded system, the computer system can also be used as part of an Internet television system, a set-top box (STB), within a printer, in a digital video disk (DVD), an industrial controller, a telephone or within an instrumentation device. Although a variety

25 of different computer systems can be used with the present invention, an exemplary general purpose computer system 112 is shown in Figure 2. PersonalJava™ can be used as a software platform of system 112.

In general, computer system 112 of Figure 2 includes an address/data bus 100 for communicating information, a central processor 101 coupled with the bus for processing information and instructions, a volatile memory 102 (e.g., random access memory RAM) coupled with the bus 100 for storing information and instructions for the central processor 101 and a non-volatile memory 103 (e.g., read only memory ROM) coupled with the bus 100 for storing static information and instructions for the processor 101. Computer system 112 also includes a data storage device 104 ("disk subsystem") such as a magnetic or optical disk and disk drive coupled with the bus 100 for storing information and instructions and a display device 105 coupled to the bus 100 for displaying information to the computer user. System 112 can also be referred to as an embedded system.

Also included in computer system 112 of Figure 2 is an optional alphanumeric input device 106 including alphanumeric and function keys coupled to the bus 100 for communicating information and command selections to the central processor 101. System 112 also includes an optional cursor control or directing device 107 coupled to the bus for communicating user input information and command selections to the central processor 101. The cursor directing device 107 can be implemented using a number of well known devices such as a mouse, a track ball, a track pad, an electronic pad and stylus, an optical tracking device, a touch screen etc. Computer system 112 can also include an optional signal generating device 108 coupled to the bus 100 for interfacing with other networked computer systems. The display device 105 utilized with the computer system 112 is optional and may be a liquid crystal

device, cathode ray tube (CRT), light emitting diode (LED), field emission device (FED, also called flat panel CRT) or other display device suitable for creating graphic images and alphanumeric characters recognizable to the user.

5 SYSTEM AND SOFTWARE ARCHITECTURES

Figure 3 illustrates a system architecture 120 that can be installed within the computer system 112 (Figure 2). The present invention includes a shared memory pool (SMP) and JavaLayer Class Manager (JCM) for allowing registered Java classes to be shared among multiple Java Virtual Machines
10 (JVMs) that simultaneously run on the system 112. By sharing basic Java classes, in lieu of maintaining a separate copy of each class for each JVM, the present invention is able to better utilize memory resources for the embedded computer system 112. A Java class, as described more fully below, contains both definitions and data and is implemented using Java byte codes which can
15 be read by a Java interpreter. Byte codes are machine independent so that they can be executed on any machine.

Within the system architecture 120 of Figure 3, the API layer 125 or application program interface is shown as the top layer. The API 125 is
20 associated with a JVM 130. Generally, only one application is resident for a particular JVM 130. The device module 135 includes a file system which can use a mini disk, a hard disk, flash ROM, a CD-ROM and/or a tape storage device. Resources that are shared 140 include hardware devices, methods, memory and Java classes that are stored in memory (e.g., the shared memory
25 pool). Communication 145 is performed using the Dispatcher, Java Native

Interface (JNI) and a Wrapper. At the bottom layers are located an operating system (OS) 150, e.g., the Aperios AV/OS, and a device driver 155.

Figure 4 is a logical block diagram of the JavaLayer software architecture 122 of the computer system in accordance with an embodiment of the present invention. At the top layer is the application 170, e.g., a HAVi application, an MHEG application, etc. Between the application 170 and the associated JVM 130 is located the interoperability API 125. At layer 175 is the event, mapper and communication portions of the JavaLayer 175. The device module can include a set-top-box 180. The operating system is shown as layer 150. The device drivers can include a board driver 182 and a serial communication driver, e.g., a 1394 board 184. Although one JVM is shown in Figure 3 and Figure 4, multiple JVMs are supported in accordance with the present invention.

15 SHARED MEMORY POOL FOR SHARING JAVA CLASSES BETWEEN MULTIPLE JVMs RUNNING SIMULTANEOUSLY

Figure 5 illustrates the JavaLayer 160 and the metaspaces with which it interworks. This architecture supports multiple JVMs simultaneously within the same computer system 112. For instance, JVMs 224 and 130 are simultaneously supported. As shown, each JVM interacts with certain Java classes 222, 226 and 228. The arrows show the direction of communication between objects and the dash-lined boxes illustrate the relationship boundaries for the application programs. When implementing a Java environment on an audio video platform, developers have to deal with the features of the JVM 130, Java classes, and audio video (AV) media applications. The JavaLayer framework supports the application

implementor in all the areas of Java and AV media applications. The framework consists of a set of constructs and rules regarding the use of the constructs. The constructs are a combination of C++ classes, Aperios active objects and Aperios meta objects. Other applications 216-220 are shown on top of the mAV
5 metaspaces 230.

Figure 6 illustrates another architecture used for memory-limited devices. Specifically, Figure 6 illustrates the JavaLayer 160 and the metaspaces with which it interworks. The architecture of Figure 6 supports multiple applications 222-228
10 (Java byte code classes) on one JVM 130 simultaneously. The arrows show the direction of communication between objects and the dash-line boxes represent the application boundaries. Also shown are applications 216-220 running on mAV metaspaces 230.

15 The JavaLayer 160 supports two architectures as shown in Figure 5 and Figure 6. JavaLayer 160 runs in the mAV metaspaces 230 in Aperios operating system v1.1. In Aperios, mDrive layer 232 is the metaspaces that support device drivers and handles interrupts and mAV metaspaces 230 is the metaspaces for AV applications. Libraries in the mAV metaspaces (which implement the JavaLayer
20 160) can communicate with drivers in the mDrive metaspaces 232 for storage, display and other device functions. In accordance with the present invention, when a class is loaded by a Java application running over the JavaLayer 160, this class is placed in common space (e.g., the shared memory pool) and is accessible by every JVM object. As shown in Figure 5, Javalayer 160 supports multiple JVMs
25 running for devices which have enough memory. Figure 6 shows the architecture

which supports multiple applications running on one JVM 130 simultaneously for memory-limited devices. In accordance with the present invention, a platform is described that allows multiple JVMs to run on a memory limited platform.

- 5 The JavaLayer framework is used to support JVM for multiple applications by sharing their basic class memory, garbage collection and a scheduler. Compared to development on a "bare" AperiOS, a more flexible and attractive development environment is provided and a shorter debugging phase is allowed mainly due to the features of the JavaLayer Framework. The JVM is primarily
- 10 PersonalJava (pJava), but it can also be any other embedded JVM written specially for the target device. JavaLayer can run multiple JVMs simultaneously on AVIOS.

- Herein, a JVM is a virtual machine which contains an execution space, an
- 15 execution engine, a class loader and a Java stack. A JVM can share classes and potentially an initialization stack with other JVM or create its own initialization stack. A JVM typically operates on Java byte codes. A created JVM can run several Java applications in parallel or serially by using the designed method. In addition, users do not need to explicitly create a thread for JVM to run. The JVMs
- 20 can potentially deploy different garbage collection and downloading mechanisms for different application. The JavaLayer framework supports a set of APIs with which applications can create a JVM, initialize a shared or non-shared class library, create a mapper for the class loader and perform event handling for special devices.

What is disclosed herein is a class management model for minimizing the amount of memory required for storing and maintaining the Java classes through class sharing and for providing mapping model to map different physical devices using their own file systems and file management APIs. Furthermore, a facility is provided which allows native applications to communicate with PersonalJava. As Figure 5 depicts, the applications 222-228 running on JVM can communicate with each other as well as with applications 216-220 running on mAV 230 by using a combination of message passing and shared memory.

With reference to Figure 7, the shared memory pool 280 is shown interacting with a JavaLayer Class Manager (JCM) 270 and a shared lock 260. When a JVM initializes its class library, the JCM 270 uses the shared lock 260 to serialize access to the shared memory pool 280, also called the class pool 280. The shared memory pool 280 is an area of memory that is established for storing Java classes and other information to be shared across multiple JVMs which are running simultaneously on system 112. In this way, the shared lock 260 is used by the present invention for serializing access of shared Java classes between multiple Java applications 256-258.

Using the shared memory pool 280 of Figure 7, a number of JVMs 251-253 can then access the same library for different operations without requiring that a separate, individual copy of the Java class be stored for each JVM as done in the prior art. In practice, the Java classes that every application needs

for operation are loaded in the shared memory pool 280 in one memory space. Every block of that memory space has a lock down (shared lock 260). If a JVM wants to access the stored Java class in the shared memory pool 280, then it must first obtain the lock associated with the desired Java class. Some of the

5 basic Java classes that can be shared between multiple JVMs are the basic Java class, for example, Java IO, Java language, and Java beans which can also have other classes included within. Some of these Java classes can consume as much as 2 Megabytes of memory each. These Java classes contain basic data required for Java. In some other programming languages,

10 apart from Java, classes are referred to as libraries.

For performance reasons physical copying and moving a class is avoided as much as possible during operations. Therefore, all classes 290-296 are managed by the central JCM 270 meta-object which manages the memory

15 pool 280 of class cells that hold all class memory. This memory pool 280 is shared between the multiple JVMs 251-253 and all JavaLayer modules as shown in Figure 7. The JCM 270 allocates memory to the classes within the shared memory pool 280 and divides the shared memory pool 280 into blocks for this purpose. The individual blocks are then managed by the JCM 270. If a

20 JVM wants to share a class, it first checks with the JCM 270 which then references a name table and the lock for the class. The JCM 270 then forwards the address of the class to the requesting JVM when the lock is free. The JVM can then access the shared Java class.

- As shown in Figure 7, an application can be associated with each JVM. As shown, application 256 is associated with JVM 251, application 257 is associated with JVM 252 and application 258 is associated with JVM 253. Requests from the JVM to access Java classes 290-296 are processed by the
- 5 JCM 270. A shared lock 260 used so that only one JVM can access any particular Java class at any time. In one embodiment, each Java class stored in the shared memory pool 280 has a read lock and a write lock. Typically, a Java class is written by only one JVM, e.g., the originator or initiator JVM, and then can be read (e.g., shared) by the other JVMs once the associated read lock is
- 10 obtained. Once a class is written into the shared memory pool 280, it becomes registered with the JCM 270 within a name table. The JCM 270 can then use the name table to respond to JVM initiated requests for access to particular Java classes.
- 15 The JavaLayer framework offers APIs to the JCM 270 for JVMs to perform operations on various classes and access the memory contents of the classes. The JavaLayer assumes a basic container, called class cells, that is used to support class loading. Class cells are dynamic memory buffers, shared by JVMs in a shared mode, and used for class manipulations avoiding physical
- 20 copying of class as much as possible. When physical drivers and application objects use the same class cells, they can share the class cell and effectively result in a JVM architecture that does not contain copies.

Figure 8 illustrates physical and logical memory organization of the

25 classes within the memory pool 280. To avoid the physical copying and

- moving of classes when operations are performed on the classes by multiple JVMs 251-253, these classes are stored in memory 310 in the manner as shown in Figure 8. The class is stored in variable sized class cells or "blocks" 290-294. The actual class is constructed by linking cells together at certain
- 5 offsets in the class cells as shown by 324a-324c. The size of any block can be dynamically adjusted based on the size of the stored Java class. The link contains the name of the class, its size and the offset from the start. The entries in the name table are shown as 322a-322c. A hash table 320 is used to reference the position within the name table for an enumerated Java class.
- 10 The name is then used to reference an offset which points to the memory position of the class within the shared memory pool 280.

- Managing the class cells as shown in Figure 8 offers the following advantageous. Class data does not have to be stored contiguously in memory
- 15 310. Also, a class cell can contain holes and therefore not all the data needs to be present. This is especially useful when reassembling PDUs. Also, cells can be shared between class cells so that copying data from one class cell to another is done by maintaining references and reference counters instead of physically copying the data. Adding data to and removing data from a class
- 20 cell can be done without physically copying or moving the class cells contents. These features lead to a zero-copy architecture, provided that the same memory management techniques are also used by the device drivers and the applications. A detailed example of the manner in which the JCM 270 (Figure 7) manipulates the class cell data is given further below.

Without class sharing of the present invention, it is difficult to invoke several JVMs for memory-limited devices such as set-top-boxes, PDS, screen phone, etc. However these devices need to support multiple applications
5 based on Java. The JavaLayer's solution is to support multi-applications on one JVM. The technique for supporting multi-applications on one JVM is to prevent the global classes or static classes from being written by more than one object. The prior art Java systems do not support the protection. The JavaLayer of the present invention provides a write-lock (WL in short) for each
10 class cell. These are maintained by the shared lock 260 of Figure 7.

As shown in Figure 9, the JCM 270 locks the class cells which are static classes or static variables to restrict the write-right. It gives the write-right to the object which is writing the static classes. JavaLayer allocates the class cell to
15 the static classes with one write-lock (WL), which is used to record the object name. The highest priority is given to the writing object to prevent dead lock. The JavaLayer guarantees the thread which holds the WL not being blocked or suspending before it releases the WL. Therefore, access to a JVM to write a Java class to the shared memory pool 280 is managed by the JCM 270 using
20 the write lock. Access to read a stored Java class from the shared memory pool 280 is managed by the JCM 270 using the shared read locks 260.

Figure 9 illustrates that any access to the shared memory pool 280 by the application 334 goes through the associated JVM 251 for that application. The

JVM 251 then accesses the JCM 270 which manages access to the shared memory pool for the Java classes 290, 292 that are stored therein.

Figure 10 illustrates steps in a process 410 used by the present invention for sharing Java classes of the shared memory pool. At step 412, a particular JVM, e.g., JVM(I), requests and obtains the write lock (e.g., "write key") (WL) to write a class, e.g., class(N), into the shared memory pool. The WL is managed by the JavaLayer class manager (JCM). When the WL is granted to JVM(I), then at step 414, it writes class(N) to the shared memory pool and the JCM registers the class name into the name table and then links the table entry with the proper offset and size in memory so that the class can be subsequently accessed. At step 416, JVM(I) then releases the write lock and class(N) is registered in the shared memory pool.

The remainder of the steps of Figure 10 illustrate the procedure used for accessing the stored class(N) that resides in the shared memory pool. At step 418, a JVM, e.g., JVM(J), checks with the JCM to determine whether or not a class(N) is stored in the shared memory pool. It is appreciated that JVM(I) and JVM(J) are both running simultaneously on system 112 (or could be the same JVM). The above determination is typically done by an instruction whose function is to establish class(N) for the JVM. During the establishment function, before the class is written into memory, a check is first done to determine if the class has already been established in the shared memory pool by another JVM. At step 420, if class(N) exists in the shared memory pool, then JVM(J) requests and obtains the read key for class(N) from the JCM. At step 422, JVM(J) then

-23-

accesses class(N), as needed, from the shared memory pool instead of copying its own instantiation of class(N) in memory. At step 424, JVM(J) then releases the read key for class(N) when done with this class. It is appreciated that steps 418-424 can be repeated for other JVMs that are simultaneously running on
5 computer system 112. Memory is thereby saved because each JVM that needs class(N) can share the instantiation of this class as established by JVM(I) without copying its own version.

PersonalJava uses a Unix-like file system for class loading. Because
10 most RT operating system platforms do not support the Unix-like system, the JavaLayer framework supports a general file system which is mapped to the real device drivers. The class mapper 175 (Figure 4) is designed to support most standard storage drivers such as remote file systems for Internet files or a Mini Disk system for local media in the SMAP (Single Media Active Project)
15 group and is analogous to a VFS layer in Unix.

Currently, the storage media is divided into three types. The first type is Local medias, such as Mini Discs, tapes, CDs. Each of the media has its own storage system. The second type is Internet-based hosts. The classes can be
20 stored in a local or remote host server which is linked to the target board by the networking. The third type are ROMs. The basic system classes and AWT classes are romized to the target board. Whatever medias the platform uses for storage of the Java classes, a standard retrieval API is necessary for the JVM to access the classes. JavaLayer framework provides a Mapper 175
25 which supports standard APIs for class management.

The native software dealing with the hardware runs on the mAV metaspace 230 (Figure 5 and Figure 6). JavaLayer communicates with these devices by using a Dispatcher class 145 (Figure 3). The standard class loader API is defined in Table I. For example, Table II illustrates an exemplary mapper that is designed for a remote file system mapping to the Unix-like file system.

Table I

```

class Unix-FileSystem {
10     FILE*fopen(const char *filename,
                const char *type);
        size_t fread(void * ptr, size_t size,
        size_t nitems, FILE *stream);
        ...;
15 }

```

Table II

```

typedef struct functab {
        char* name;
20     void* ftnlc;
    } flinctab;

    flinctab Mapper = {
        { "fopen 15Unix_FileSystemFIPvUiPui", rfs_open},
25     { "fopen_15Unix_FileSystemFIPvUiPui", miniDisk_open},
        { "fopen_15Unix_FileSystemFIPvUiPui", flashRom_open},
        ...
    }

```

30 A mapper generator is designed for the automatic mapper generation by setting JVM class loader type.

COMMUNICATION MANAGEMENT USING THE SHARED MEMORY POOL

Communication management using mAV's communication mechanisms is now discussed. The shared memory pool 280 can be used for message passing between applications. For instance, a native application can use the JavaLayer to communicate with a Java application. The JavaLayer framework supports communication between Java execution spaces and native applications running on AVIOS or some other metaspace in Aperios using the Dispatcher 145. The Dispatcher class 145 organizes programs as a collection of actions, each of which is invoked upon receipt of a message. A program registers a set of methods which the dispatcher calls in the context of its threads when a message arrives. In addition, a Java Dispatcher class 145 can be extended to support Java and can be used for AV applications. For instance, JVM can support an API that accepts the Dispatcher 145 from the Aperios operating system 150 and invokes itself so that it can run in the application (JVM). The API interfaces between the JVM and the operating system 150 (Figure 4).

Refer to Figure 11. Within the communication method of the present invention, real time coordination is allowed both between JVM and native applications and between JVMs via the shared memory pool 280 for an embedded system. Figure 11 illustrates an example system 510 for coordination between a JVM 130 and native applications 516. In this example, Home Audio/Video Interoperability (HAVi) application 516 needs to download a Device Control Module (DCM) 512 to install a device and control it. HAVi

516 first runs an application 514 and this application 514 needs a new device. Therefore, HAVi 516 downloads the device DCM 512, as shown by 520, and stores it in the shared memory pool 280 of the present invention via interface 526. After that, HAVi 516 informs the JVM 130, as shown by 524, to run the DCM module to install the new device. The JVM 130 then accesses the module that is stored in the shared memory pool 280 via 528.

In the above fashion, the shared memory pool 280 is providing a message communication pathway between the native application and the JVM. Typically, the message from the native application to the Java application is processed using help from the operating system's (e.g., Aperios) message passing abilities. A target code in the message can be used to identify whether the message is for a Java application or for a native application. If the message contains a Java code, then the operating system will communicate with the JVM.

HAVi 516 of Figure 11, the shared memory pool 280 and the JavaLayer 160 interface with the operating system 150 and the JVM 130 interfaces with the JavaLayer 160. Havi 516 can access or destroy contents of the shared memory pool 280 via interface 526. The JVM 130 can run and access information of the shared memory pool 280 via interface 528. JavaLayer 160 provides management via the JCM 270 as represented by interface 530 in Figure 11. Associated with JavaLayer 160, there is the shared memory pool (SMP) manager, also called JCM 270 and a JVM manager.

In accordance with this example, the JCM 270 provides the following services as shown in Table III below.

5 Table III

Service Value	Parameters	Return
10 SMPManager::GetSMP	(size)	MemoryID
SMPManager::GetSMPFreeSize		size
15 SMPManager::GetSMPStatus	(MemoryID, Offset)	RUN, COMPLETE, FAIL
SMPManager::GetSMPDestroy	(MemoryID)	FAIL, SUCCESS

20

The JVM manager provides the following services as shown in Table IV below:

25 Table IV

Service	Parameters	Return Value
30 JVMManager::Execute	(ModuleName, MemoryID, Offset),	FAIL, SUCCESS
JVMManager::SetStatus	(MemoryID, Offset, Status)	FAIL SUCCESS

35 Figure 12 and Figure 13 illustrate an example of coordination between several JVMs using the shared memory pool 280 of the present invention. With reference to Figure 12, a first JVM 130a and a second JVM 130b are shown interfacing with the JavaLayer 160. The first JVM 130a is receiving a video

stream from a broadcasting source 546 and the second JVM 130b is playing the video stream simultaneously 548 to some rendering device (e.g., a DTV). The first and second JVMs are communicating the video stream 544 using the shared memory pool 280, with the first JVM 130a writing into the shared

5 memory pool 280 and the second JVM 130b reading from the shared memory pool 280. In home networking, the first JVM (called JVM1) 130a is running an application 546 for receiving the video stream from a broadcast and decoding the stream in the shared memory pool 280. The second JVM (called JVM2) 130b was receiving user's control and playing the stream.

10

Figure 13 illustrates a timing diagram of events occurring between the coordination of the first JVM 130a (on the left) and the second JVM 130b (on the right). At 562, the first JVM 130a receives the video stream and by 574 it saves this information into the shared memory pool (SMP) 280. The second

15 JVM 130b, under user control, decodes and plays back this video at 568 using interface 576. At 564, the first JVM 130a receives a new decoding module that is downloaded and saved to the SMP 280. The second JVM 130b then decodes and plays the new module as shown by 580. At 584, the second JVM 130b sends a command to play a game through the Internet as instructed by

20 user control 570. At 566, the first JVM 130a receives and sends back a stream 590 in response to the Internet game. At interface 586, information is both saved and retrieved by the first JVM 130a. The second JVM 130b then supports game playing 572 by interfacing with the SMP 280 to provide interoperability 588. In all cases of Figure 13, the shared memory pool 280 is

25 used as an intermediary in the communication method.

Event Handling is described. DeviceEvent is the base interface for events in the JavaLayer framework. In order to support the Java AWT, event model, an implementation of DeviceEvent, is required to sub-class

5 java.util.EventObject. Device events asynchronously deliver information about keyboard input, screen touch, remote control. To receive events, an object implements the device listener interface and uses the "add device event" method to the event queue. JavaLayer provides an deviceEventManager to map different devices to the standard event handling interfaces. The definition

10 is shown in the list of Table V.

Table V

```

classAV_EventHandling {
    Position      x, y;
15   EventType    evType;
    Event         *queue;
    void (*deviceEventManager(DeviceID dID)) (void);
    ...;
20 }

typedef struct functab {
    DeviceID      dID;
    void*         func;
25 } Device_Func_tab;

Device_Func_tab DeviceMapper = {
    {Keyboard, KeyboardHandling},
    {RemoteController, RemoteControllerHandling},
30   {TouchPanel, TouchPanelHandling},
    ...
}

```

An example implementation of a JVM within the AV/OS operating

35 system is presented below. The syntax of the list of Table VI is described as below.

- Regarding the syntax address server, the JVM uses the host of "serverHost" as a server to download classes. The target devices connect with the host by the Internet. Regarding the syntax version ver, this records the WM version information. Regarding the syntax stack stackentry, this records the JVM stack entry. Regarding the syntax JVM *share, this is the shared JVM point. If it is not NULL, the stackentry is NULL. The size_t heap is the heap size for this JVM. Here, the "ClassLoader" is shared.
- 10 The syntax static JVM *Create (size_t heap, Version ver, uint32 threadopts, JVM *share) constructs a new JVM with a new running thread and the heap size. It implicitly initializes basic Java Classes, with pJava version information and also checks whether the version is compatible with the implemented JVM or not. With the "share" parameter, the WM does not
- 15 initialize Java Classes but shares Class Stack with another JVM pointed by param share. Two JVMs, which are expected to share the classes stack, should be compatible versions and synchronized. This parameter can be pointed only at the creating time. Otherwise, the JVM owns an independent Class Stack by calling to initialisejavaClass (uint32 level). Users initialize
- 20 Class Stack explicitly from ROM / RAM / Flash memory. This method also can change the current JVM to a not-sharing JVM from a sharing status. Upon successful completion, the pointer to the new JVM is returned. Otherwise NULL is return.

The syntax `initializeJVMclass(class, method)` is used to establish a class for the JVM. If the class already exists in the shared memory pool 280, then this command does not generate another a copy of the class within the shared memory pool 280 but rather uses the copy that already exists there. If the class

5 does not exist within the shared memory pool, then this command generates the copy of the class into the shared memory pool and registers the class with the JCM 270. The command `executeJavaClass()` is used for message sending.

- 10 Regarding the syntax, status `Destroy()`, it is used for destruction, all activities in the JVM are explicitly terminated and destroyed. The initialized class stack also is freed. If there are any other JVM is sharing class stack with the current JVM, the destruction is pending until the release of the class stack. The syntax `* void executejavaClass (char * className, void * args[])`
- 15 represents the method for notation of running "application class." For example, the JVM runs the application bytecode class "HelloWorld.class", prints out "Hello World" if it is used as:

```
executeJavaClass ("HelloWorld", "main")
```

20 Table VI

```

class JVM {
    static JVM * Creat();
    static JVM * Creat(size_t heap, Version ver,
25     JVM *share, Address server,
        Class Loader cl_type);
    Version ver;
    Stack stackentry;
    JVM * share;
    size_t heap;
30     enum ClassLoader {
        MD,

```

-32-

```

        rfs_FILE,
        ROM,
        Download
    } cl_type;
5   Address server;
    initializeJVMclass(class, method);
    executeJavaClass();
    status Destroy();
10  void SetServer (Address addr);
    Stack initialiseJavaClass (uint32 level);
    AV_EventHandling *event;
    Unix_FileSystem fileHandling;
}

```

15 An implementation of PersonalJava (pJava) is discussed. A combination example of the modules of JavaLayer is pJava. It is appreciated that pJava is a Java™ API and Java Application Environment for networked applications running on personal consumer devices. Running pJava on Aperios is targeted for platforms such as set-top boxes which are composed of

20 Aperios real time operating system, download run time environment, Internet-based web browser and home network.

The list of Table VII shows an example of pJava which is created with heap size 2097144, PersonalJava 1.1 compatible, and running an application

25 called "Browser".

Table VII

```

pJava = JVM::Creat(2097144, pJaval.1,
30  NULL, 0, rfs_FILE);
pJava -> SetServer (43.135.113.83);
pJava-> initializeJVMclass();
pJava-> executeJavaClass(Browser, "main");
pJava-> executeJavaClass("SmapApplet.class", "run");

```

35 The preferred embodiment of the present invention, a computer-implemented method and system for sharing Java classes across multiple Java

virtual machines on the same computer system and a method for communication between Java and native applications (or between two or more JVMs), is thus described. While the present invention has been described in particular embodiments, it should be appreciated that the present invention

5 should not be construed as limited by such embodiments, but rather construed according to the below claims.

CLAIMS

What is claimed is:

- 5 1. In a computer system, a method
comprising the steps of:
- a) a first Java Virtual Machine (JVM) loading an instantiation of a first
Java class into a shared memory pool;
 - b) registering a name of said first Java class into a name table that is
10 used for said shared memory pool;
 - c) querying said name table on behalf of a second JVM to determine if
said first Java class is stored in said shared memory pool, wherein said second
JVM is operating on said computer system simultaneously with said first JVM;
and
 - 15 d) said second JVM accessing and utilizing said instantiation of said first
Java class from said shared memory pool in lieu of creating a separate
instantiation of said first Java class for its own use.
2. A method as described in Claim 1 further comprising the steps of:
- 20 said second JVM loading an instantiation of a second Java class into
said shared memory pool;
- registering a name of said second Java class into said name table;
 - querying said name table on behalf of a third JVM to determine if said
second Java class is stored in said shared memory pool, wherein said third
25 JVM is operating on said computer system simultaneously with said first and
said second JVMs; and

said third JVM accessing and utilizing said instantiation of said second Java class from said shared memory pool in lieu of creating a separate instantiation of said second Java class for its own use.

3. A method as described in Claim 1 wherein said first JVM has an associated first Java application and wherein said second JVM has an associated second Java application.

4. A method as described in Claim 1 wherein said first Java class is a Java IO class.

5. The method of any one of Claims 1-4 wherein the computer system is an embedded computer system having a processor coupled to a bus and a memory device coupled to said bus wherein said memory device contains instructions stored therein that implement said method being a method of sharing Java classes.

6. The method of Claim 1, wherein the method is a method of communicating between applications further comprising the step of:

e) a first Java application of said first JVM communicating messages to a native application by writing said messages into said shared memory pool and said native application reading said messages from said shared memory pool.

7. The method as described in Claim 6 further comprising the step of said native application communicating messages with said first Java application by writing said messages into said shared memory pool and said Java application reading said messages from said shared memory pool.

8. The method as described in Claim 6 further comprising the step of said first Java application communicating messages with a second Java application of said second JVM by said first Java application writing said messages into said shared memory pool and said

second Java application reading said messages from said shared memory pool.

9. The method as described in Claim 1, 5 or 6 wherein said step a) comprises the steps of:

- a1) accessing a write lock associated with said first Java class;
- a2) granting write access of said first Java class to said first JVM when said write lock associated with said first Java class becomes free; and
- a3) said first JVM writing said instantiation of said first Java class in response to said step a2).

10. The method as described in Claim 9 wherein said step a3) comprises the steps of:

loading said instantiation of said first Java class into blocks of said shared memory space and registering offset and size information regarding said blocks; and
linking said offset and size information with said name of said first Java class.

11. The method as described in Claim 1, 5 or 6 wherein said step d) comprises the steps of:

- d1) accessing a read lock associated with said first Java class;
- d2) granting access of said first Java class to said second JVM when said read lock associated with said first Java class becomes free; and
- d3) said second JVM accessing said instantiation of said first Java class in response to said step d2).




FIGURE 1




FIGURE 2

120




FIGURE 3




FIGURE 4

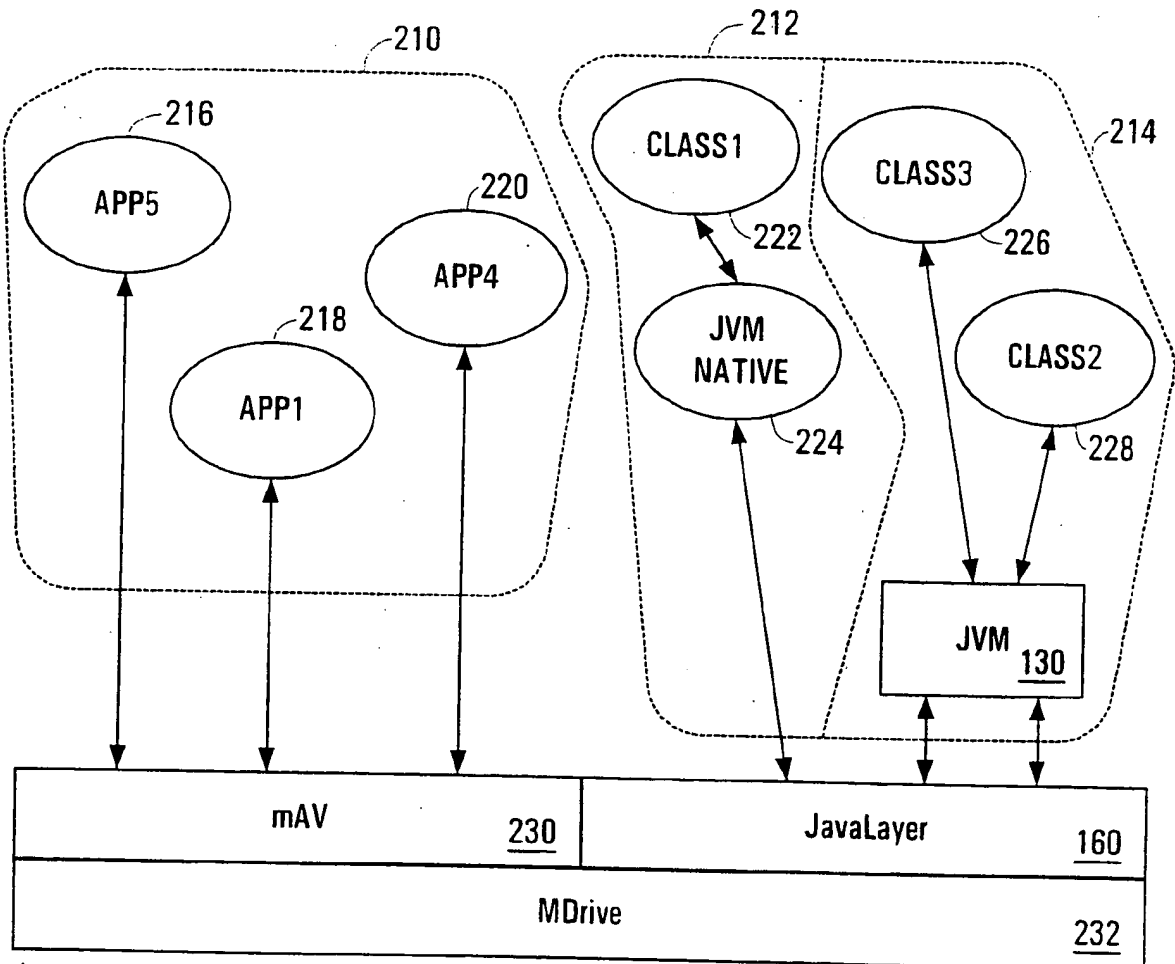


FIGURE 5

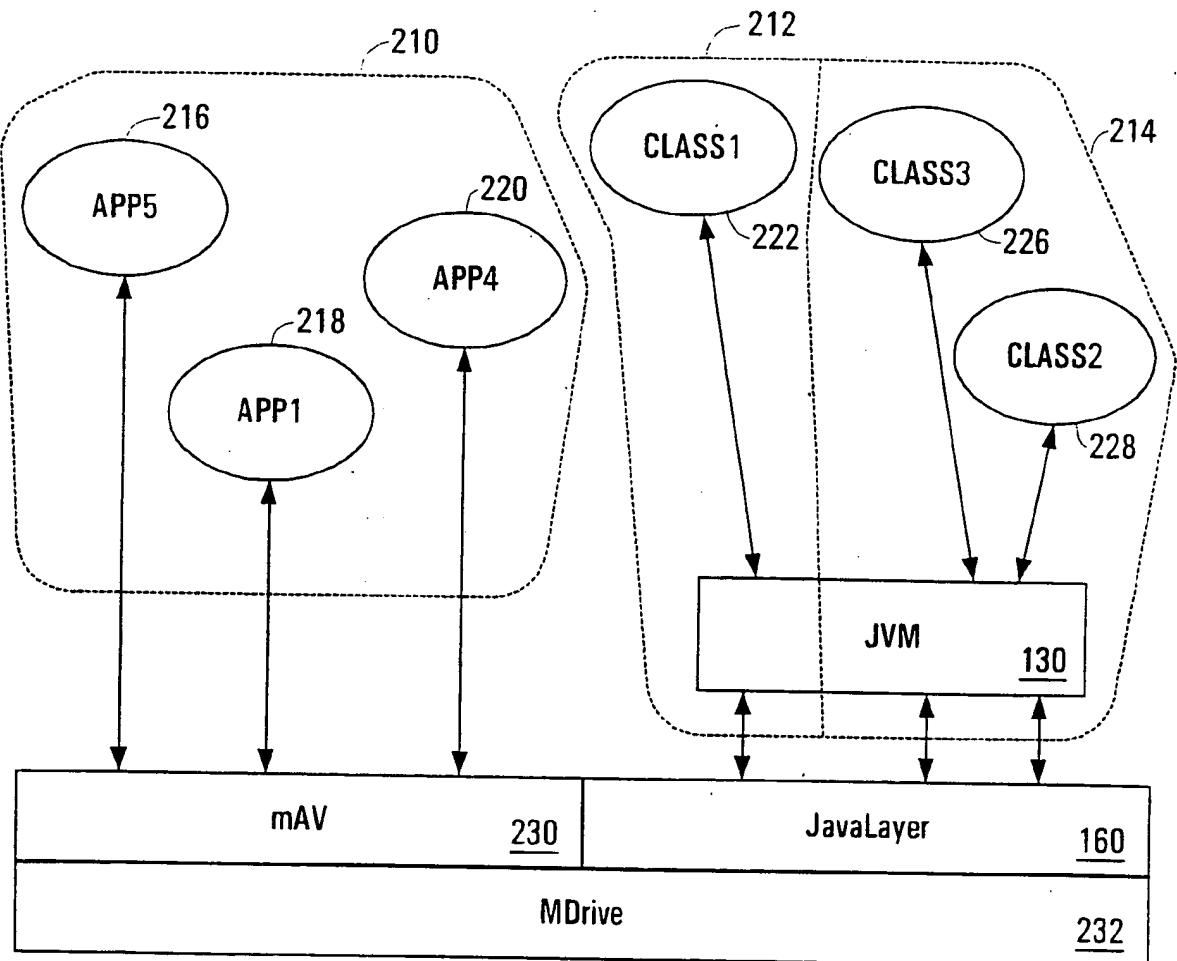


FIGURE 6




FIGURE 7




FIGURE 8




FIGURE 9

410

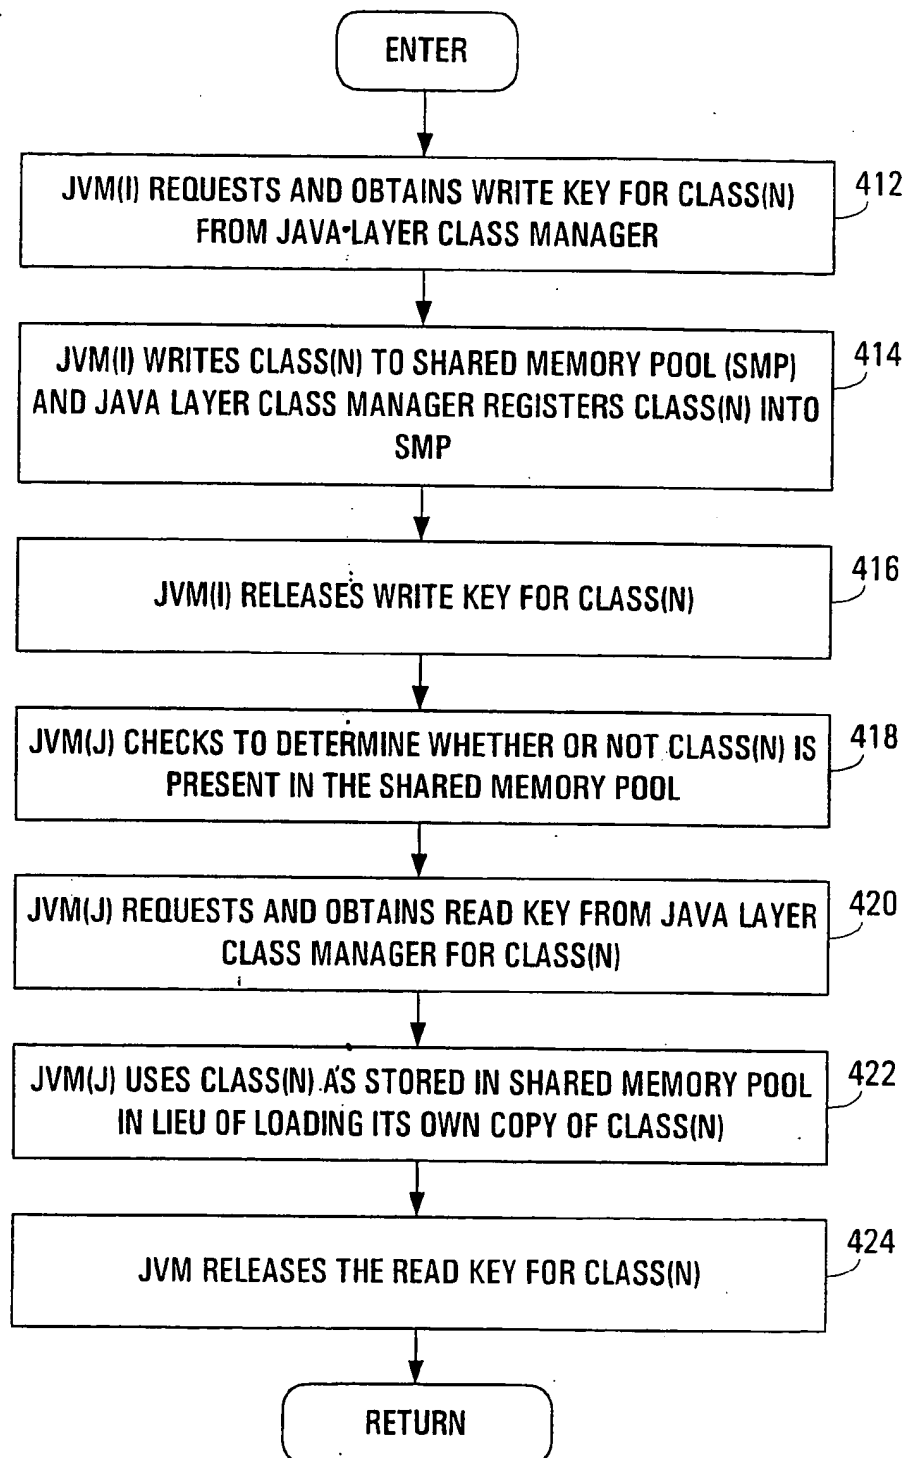
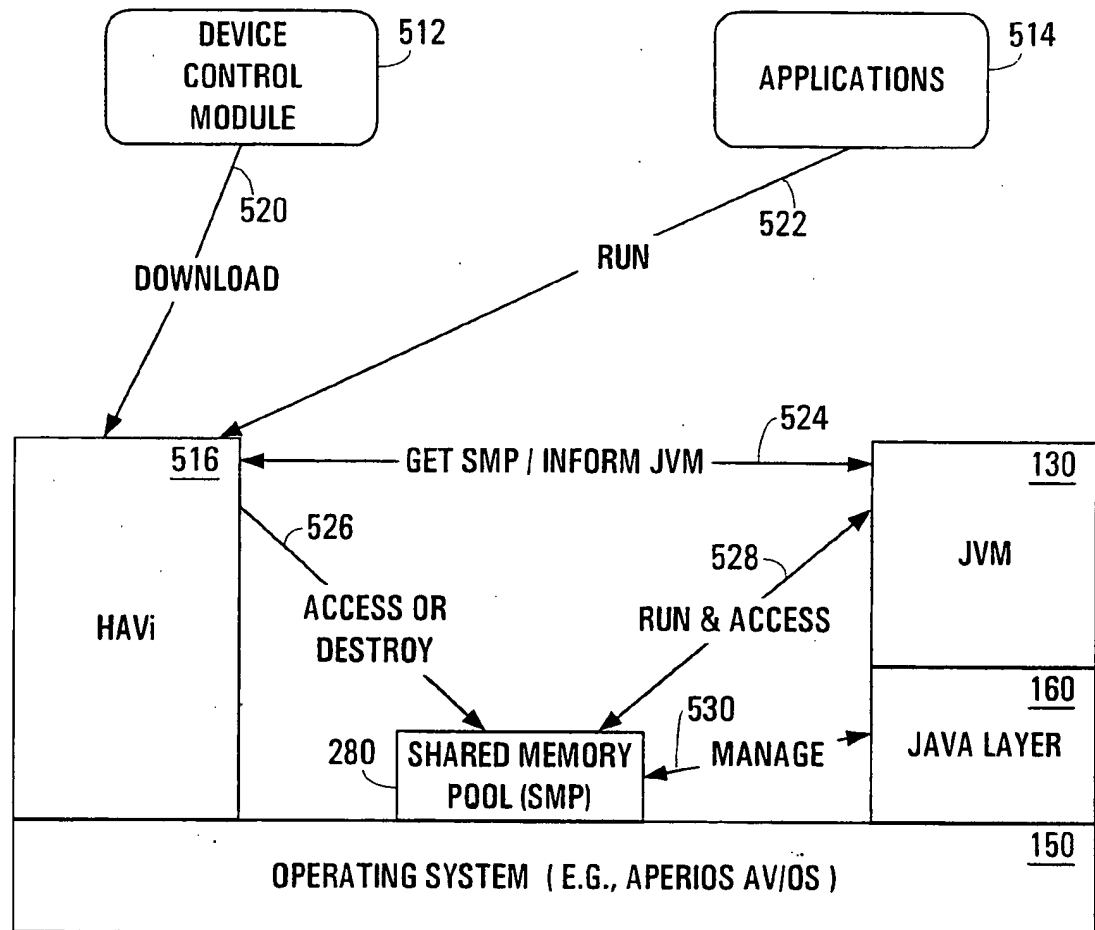


FIGURE 10

510**FIGURE 11**

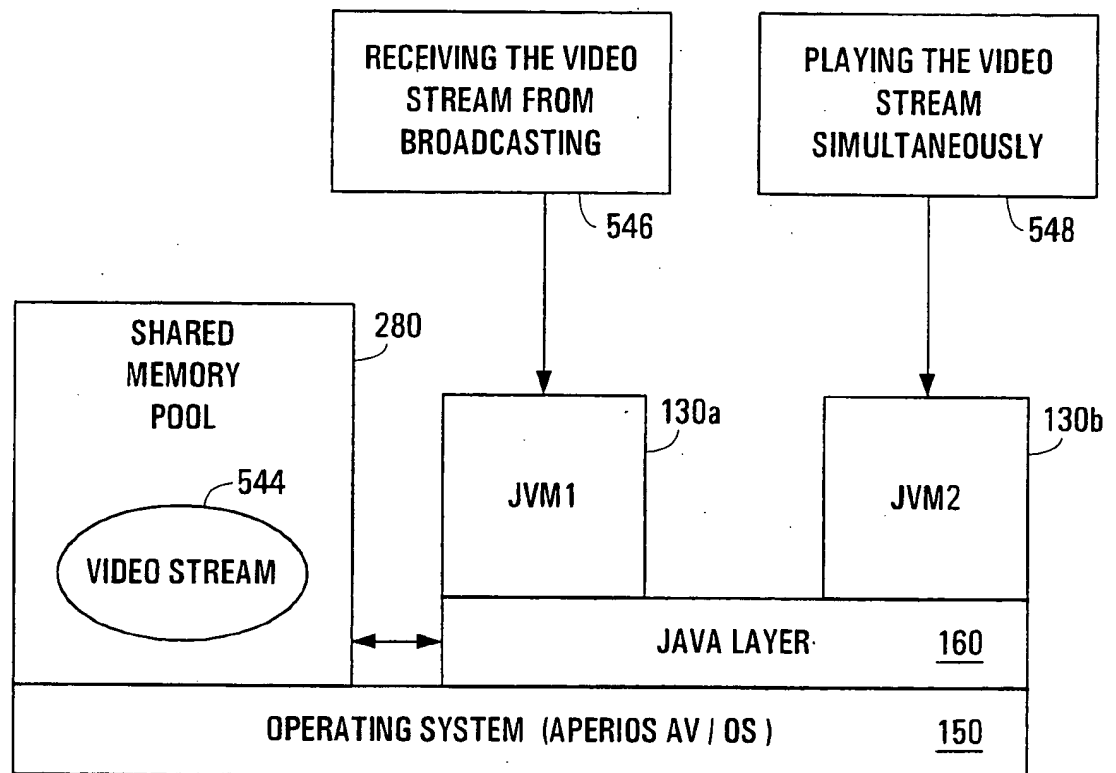
540

FIGURE 12

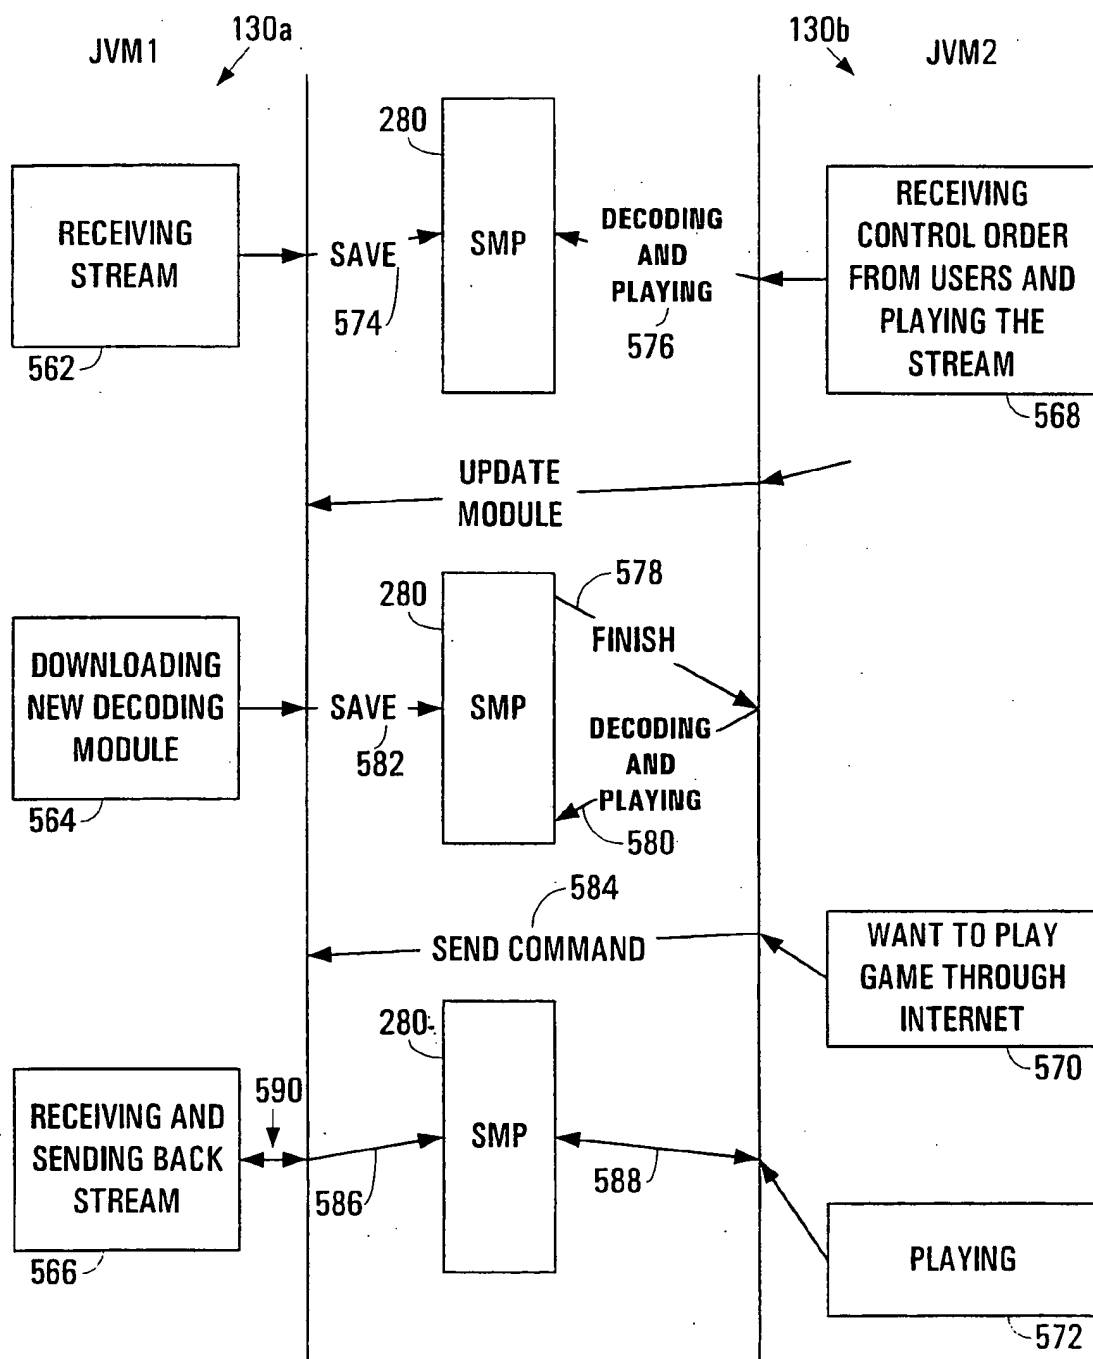
560

FIGURE 13

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 99/28859

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F9/445

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, IBM-TDB, INSPEC, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 822 580 A (LEUNG WYATT) 13 October 1998 (1998-10-13) column 20, line 1 -column 21, line 12 claims 1,2	1-3,5
A	VENNERS, BILL: "Inside the Java Virtual Machine (page 78)" December 1997 (1997-12), MCGRAW-HILL, US XP002144225 the whole document	1-5,9-11
A	EP 0 755 004 A (NIPPON ELECTRIC CO) 22 January 1997 (1997-01-22) page 13, line 44 - line 50	1-3,5
A	US 5 742 813 A (BEALL CHRISTOPHER W ET AL) 21 April 1998 (1998-04-21) column 8, line 65 -column 10, line 55	1,9,11
	-/--	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

7 August 2000

Date of mailing of the international search report

18/08/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Bijn, K

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 99/28859

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 737 605 A (CUNNINGHAM CONNEL G ET AL) 7 April 1998 (1998-04-07) column 3, line 40 -column 6, line 3 -----	1-3,5
P,A	DAVIS, THOMAS E.: "Use JNDI to share objects between different virtual machines" JAVAWORLD, 'Online! July 1999 (1999-07), pages 1-9, XP002142437 Retrieved from the Internet: <URL:http://www.javaworld.com/javaworld/jw-07-1999/jw-07-cooltools.html> 'retrieved on 2000-07-13! the whole document -----	1-3,5

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 99/28859

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5822580 A	13-10-1998	CA 2238973 A DE 69701623 D EP 0861467 A WO 9726597 A	24-07-1997 11-05-2000 02-09-1998 24-07-1997
EP 0755004 A	22-01-1997	US 5745703 A JP 3006675 B JP 9114791 A	28-04-1998 07-02-2000 02-05-1997
US 5742813 A	21-04-1998	US 5838965 A AU 4462596 A CA 2204733 A EP 0803102 A WO 9615501 A US 5778356 A US 5835910 A	17-11-1998 06-06-1996 23-05-1996 29-10-1997 23-05-1996 07-07-1998 10-11-1998
US 5737605 A	07-04-1998	DE 69418963 D EP 0647902 A JP 7121373 A	15-07-1999 12-04-1995 12-05-1995

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☒ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.